

I noticed a few nit-picky items that need to be clarified or changed on the Energy Star 4.0 spec before the final publication. Since servers have several exceptions, extra care needs to be taken to keep from writing blanket statements that don't apply to servers. While most servers that interest Hewlett-Packard are not covered by this specification, I would personally like to see Tier II defined such that any well-designed HP server can be designed to be Energy Star compliant, regardless of how big it is.

In the Draft 1 spec, I made the following observations:

- Page 6, lines 213-215. The definition of a "desktop-derived server" seems to be more like that of an enterprise server (or in your terminology "Mid-range or Large server"). I don't view "desktop-derived servers" to be designed for mission critical applications nor to be utilized 24 hours a day nor to be available with only minutes of downtime per year. That level of functionality and availability is more of an enterprise server definition. So, perhaps the bullet item on lines 213-215 should be moved under the lines 223-236 definition for mid-range or large servers. Desktop-derived servers by definition don't have high-availability, since they do not have the redundant systems of an enterprise-class server. For consistency, the table on page 9, line 349 should say "mid-range or large servers" instead of "mid- to large-servers".
- Page 8 needs to have the term "Power Management" defined. On page 14, they write about how that Energy Star 4.0 computers must be shipped with Power Management turned on, and what must happen in a power managed system. I think that when the spec talks about "power management", you are thinking of the Windows control panel type of selections where it defines when to turn off monitors and hard drives, and when to go to stand-by and hibernate modes. Servers are not intended to sleep, and on page 11 the Energy Star 4.0 spec says that Sleep Mode is N/A. So servers will not be shipped with "power management" turned on, and there can be no blanket statements about all computers being shipped with power management turned on. HP would define "power management" for servers to have a completely different meaning, so perhaps the best way to avoid confusion would be to make "power management" not required for servers of any kind. Server power reductions should come from designed-in energy efficiency at full performance.
- Page 15 has some typos. The word "remain" is misspelled on line 477. The word "effective" is misspelled on line 479. The table under line 485 mentions "Version 2.0" and I can't tell if that should be "Version 4.0" or if that is referencing some other specification. It's unclear about what specification it is talking about, regardless.
- Page 17 typos: Lines 529 and 536 use the terms "Tier 1" and "Tier 2", whereas everywhere else those stages are listed in Roman numerals (e.g. Tier I and Tier II).
- I don't have an Idle State power level recommendation for desktop-derived servers at this time. This is a complex subject that needs to take into consideration what hardware is installed. Server hardware is optimized to the applications that run on them, and we can't treat them like cookie-cutter PCs. We believe that (if defined) servers are better served in this parameter by defining a minimum power reduction achieved during idle (versus the power at high utilization). This will reward systems that (for instance) ship with ACPI C1e states enabled in their operating system and optimized in microprocessors to yield as low power as possible when the operating system is idle and the processor has executed a Halt instruction.

Warm Regards,

Roger Tipley
Hewlett-Packard Company
Office of the Chief Technology Officer
Industry Standard Server group
Houston, TX